

WHAT IS CLAIMED IS:

1. In a system having a plurality of storage elements, a method for maintaining stored objects in the plurality of storage elements, comprising the steps of:

5 storing an object in a plurality of storage elements; in response to a request to update the object, determining a future invalidation time when the object should be invalidated in storage elements currently storing the object; and

10 delaying updating the object until the invalidation time has passed.

2. The method as recited in claim 1, further comprising a step of serving an object by a storage element only if the invalidation time for the object has not passed.

15

3. The method as recited in claim 1, wherein the system includes a consistency coordinator and the method further comprises steps of:

20

determining the invalidation time using the

consistency coordinator; and

communicating the invalidation time to the storage elements storing the object.

5 4. The method as recited in claim 1, further comprising a step of sending, from the consistency coordinator to a storage element and/or from a storage element to the consistency coordinator, heart beat messages to obtain availability information.

10

 5. The method as recited in claim 4, further comprising declaring an entity expecting a heartbeat down in response to the entity failing to receive a heart beat.

15

 6. The method as recited in claim 1, further comprising a step of sending an acknowledge message, by the storage elements, that the invalidation time has been received.

20

 7. The method as recited in claim 1, further comprising a step of invalidating all objects that have a current time later than the invalidation time.

8. The method as recited in claim 1, further comprising a step of discarding and replacing a current version of an object with a pending-update version if the
5 object has a current time later than the invalidation time.

9. The method as recited in claim 1, wherein the storage elements include at least one cache.

10 10. A program storage device readable by machine, tangibly embodying a program of instructions executable by the machine to perform method steps for a method for maintaining stored objects in a plurality of storage elements, the method steps comprising:
15 storing an object in a plurality of storage elements; in response to a request to update the object, determining a future invalidation time when the object should be invalidated in storage elements currently storing the object; and
20 delaying updating the object until the invalidation time has passed.

11. A method for maintaining stored objects
comprising the steps of:

providing a plurality of storage elements;

storing an object in a plurality of storage elements;

5 requesting an update for the object;

determining a future invalidation time when the object
is to be invalidated in storage elements currently storing
the object; and

delaying updating the object until the invalidation
10 time has passed.

12. The method as recited in claim 11, further
comprising a step of serving an object by a storage element
only if the invalidation time for the object has not
15 passed.

13. The method as recited in claim 11, further
comprising steps of:

determining the invalidation time using a consistency
20 coordinator; and

communicating the invalidation time to the storage
elements storing the object.

14. The method as recited in claim 13, further comprising a step of sending heart beat messages to obtain availability information from the consistency coordinator to a storage element and/or from a storage element to the consistency coordinator.

15. The method as recited in claim 14, further comprising declaring an entity expecting a heartbeat down in response to the entity failing to receive a heart beat.

16. The method as recited in claim 11, further comprising a step of sending an acknowledge message, by the storage elements that the invalidation time has been received.

17. The method as recited in claim 11, further comprising a step of invalidating all objects that have a current time later than the invalidation time.

18. The method as recited in claim 11, further comprising a step of discarding and replacing a current

version of an object with a pending-update version if the object has a current time later than the invalidation time.

19. The method as recited in claim 11, wherein the
5 storage elements include at least one cache.

20. A program storage device readable by machine,
tangibly embodying a program of instructions executable by
the machine to perform method steps for method for
10 maintaining stored objects, the method steps comprising:
providing a plurality of storage elements;
storing an object in a plurality of storage elements;
requesting an update for the object;
determining a future invalidation time when the object
15 is to be invalidated in storage elements currently storing
the object; and
delaying updating the object until the invalidation
time has passed.

20 21. A system for maintaining stored objects
consistently, comprising:
at least one object;

a plurality of storage elements for storing the object; and

a consistency coordinator which receives a request to update the object from at least one the storage elements and in response determines a future invalidation time when the object should be invalidated in storage elements currently storing the object;

the storage elements delaying updating the object until the invalidation time has passed.

10

22. The system as recited in claim 21, wherein the storage elements serve an object only if the invalidation time for the object has not passed.

15

23. The system as recited in claim 21, wherein the consistency coordinator communicates the invalidation time only to the storage elements storing the object.

20

24. The system as recited in claim 21, further comprising heart beat messages sent to obtain availability information from the consistency coordinator to a storage element and/or from a storage element to the consistency

coordinator.

25. The system as recited in claim 24, further comprising a down mode for which an entity expecting a heartbeat declares in response to the entity failing to receive a heart beat.

26. The system as recited in claim 21, further comprising an acknowledge message sent by the storage elements when the invalidation time has been received.

27. The system as recited in claim 21, further comprising a pending-update version list to which an object is added pending an update if the object has a current time later than the invalidation time.

28. The method as recited in claim 21, wherein the storage elements include at least one cache.

20